CLAIMS

What is claimed is:

1. In a communication network, a method of TCP state migration comprising the steps of:

a) establishing a communication session between a client and a front-end node at a first bottom TCP (BTCP) module located below a first TCP module in a first operating system at said front-end node, said front-end node accessing a plurality of back-end web servers forming a web server cluster that contains content;

b) receiving a HTTP request from said client at said first BTCP module;

c) parsing said HTTP request to determine which back-end web server, a selected back-end web server, in said plurality of back-end web servers can process said HTTP request, said selected back-end web server not said front-end node;

d) extending said communication session to said selected back-end web server by handing-off an initial TCP state of said first BTCP module to said selected back-end web server;

e) sending said HTTP request to said selected back-end web server;

f) switching a bottom IP (BIP) module at said front-end node to a forwarding mode, wherein packets received at said BIP module from said client are forwarded to said

selected back-end web server, said BIP module located below an IP module at said front-end node; and

g) terminating said communication session at said front-end node after said HTTP request is fully processed.

2. The method as described in Claim 1, wherein said content is partially replicated between each of said plurality of back-end web servers.

3. The method as described in Claim 1, wherein said second BTCP module is located below a second TCP module in a second operating system at said selected back-end web server.

4. The method as described in Claim 1, wherein said initial TCP state is associated with said communication session, said communication session established for the transfer of data contained within said content to said client.

5. The method as described in Claim 4, wherein said step d) comprises the further steps of:

sending a SYN packet to said selected back-end web server,

said SYN packet intercepted by said second BTCP module, said SYN packet originally sent from said client to said front-end node in requesting said communication session, said SYN packet stored at said first BTCP module;

including an initial sequence number within said SYN packet that enables said second BTCP module to understand a proper TCP state of said first BTCP module in said communication session;

receiving a SYN/ACK packet from said selected back-end web server, said SYN/ACK packet updated by said second BTCP module to reflect said proper TCP state of said first BTCP module; and

sending an ACK packet from said first BTCP module to said selected back-end web server, said ACK packet originally sent from said client to said front-end node in establishing said communication session.

6. The method as described in Claim 1, wherein said method comprises the further step of:

sending response packets from said selected back-end web server to said client in a communication path that does not include said front-end node by changing headers of said response packets such that it appears that the source of said response packets is said first BTCP in its proper TCP state.

7. The method as described in Claim 1, wherein step g) comprises the further steps of:

intercepting TCP control packets from a second TCP module located at said selected back-end web server at said second BTCP module;

sending said TCP control packets to said first BTCP module from said second BTCP module;

sending said TCP control packets to said client from said first BTCP module; and

terminating said communication session at said front-end node and said back-end web server.

8. The method as described in Claim 1, wherein said front-end node and said plurality of back-end web servers comprise a web site, said front-end node providing a virtual IP address for said web site.

9. The method as described in claim 8, wherein said front-end node, and said plurality of back-end web servers are coupled together by a local area network.

10. The method as described in Claim 8, wherein said front-end node and said plurality of back-end web servers are coupled together by a wide area network.

11. In a communication network, a method of TCP state migration comprising the steps of:

a) receiving a request from a client for establishing a communication session at a first bottom TCP (BTCP) module located at a front-end node, said front-end node accessing a plurality of back-end web servers containing content, wherein said content is partially replicated between each of said plurality of back-end web servers, said communication session established for the transfer of data contained within said content to said client;

b) establishing said communication session between said client and said first BTCP module, said first BTCP module located below a first TCP module in a first operating system at said front-end node;

c) receiving a HTTP request from said client at said first BTCP module;

d) parsing said HTTP request to determine which back-end web server, a selected back-end web server, in said plurality of back-end web servers contains said data in order to process said HTTP request, said selected back-end web server not said front-end node;

e) extending said communication session to said selected back-end web server by handing-off an initial TCP state of said first BTCP module to a second BTCP module located at said selected back-end web server, said initial TCP state associated with said communication session between said client and said first BTCP module, said second BTCP module located below a second TCP module in a second operating system at said selected back-end web server;

f) sending said HTTP request to said selected back-end web server;

g) switching a bottom IP (BIP) module in said front-end node to a forwarding mode, wherein packets, from said client, received at said front-end node are intercepted by said BIP module and forwarded to said selected back-end web server, said BIP module located below an IP module in said front-end node, said BIP module changing destination IP addresses of said packets to said selected back-end web server; and

h) terminating said communication session after said HTTP request has been fully processed.

12. The method as described in Claim 1, wherein step e) comprises the further steps of:

e1) storing a SYN packet sent from said client to said front-end node, said SYN packet requesting said communication session in step a);

e2) storing an ACK packet sent from said client to said front end node in establishing said communication session;

e3) sending said SYN packet to said selected back-end web server so that it appears that said SYN packet originated from said client;

e4) sending said initial TCP state to said second BTCP module, including said initial sequence number, that enables said second BTCP module to understand a proper TCP state of said first BTCP module for said communication session;

e5) receiving a SYN/ACK packet at said first BTCP module from said second TCP module, said SYN/ACK packet updated by said second BTCP module to reflect said proper TCP state at said first BTCP for said communication session; and

e6) sending said ACK packet to said selected back-end web server to extend said communication session to said selected server.

13. The method as described in Claim 12, wherein step e4) includes the further step of including said initial sequence number in said SYN packet.

14. The method as described in Claim 11, wherein said method comprises the further step of sending response packets from said back-end web server to said client in a communication path that does not include said front-end node, by changing headers of said response packets such that it appears that the source of said response packets is said front-end node with said proper TCP state.

15. The method as described in Claim 11, wherein step h) comprises the steps of:

intercepting TCP control packets from said selected back-end web server at said second BTCP module;

sending said TCP control packets to said first BTCP module from said second BTCP module;

sending said TCP control packets to said client from said first BTCP module; and

terminating said communication session at said front-end node and said back-end web server.

16. The method as described in Claim 15, wherein said TCP control packets include a RST flag and a FIN flag.

17. The method as described in Claim 11, wherein said method bypasses the TCP module at said front end server.

18. The method as described in Claim 11, wherein said front-end node, and said plurality of back-end web servers comprise a web site, said front-end node providing a virtual IP address for said web site.

19. The method as described in claim 18, wherein said front-end node, and said plurality of back-end web servers are coupled together by a local area network.

20. The method as described in Claim 18, wherein said front-end node and said plurality of back-end web servers are coupled together by a wide area network.

21. The method as described in Claim 11, wherein said content is partitioned between each of said plurality of back-end web servers.

22. A communication network for TCP state migration comprising:

a client;

a front end server coupled to said client by said communication network, said front-end node including a front-end bottom TCP (BTCP) module located below a front-end TCP module in a first operating system, and a bottom IP (BIP) module located below an IP module in said first operating system; and

a plurality of back-end web servers including a selected back-end web server, said plurality of back-end

web servers containing content that is partitioned between each of said plurality of back-end web servers, each of said plurality of back-end web servers coupled to said front-end node through said communication network, each of said plurality of back-end web servers including a back-end bottom TCP module located below a back-end TCP module.

23. The communication network as described in Claim 22, wherein said front-end BTCP module establishes a communication session with said client for the transfer of data contained within said content to said client.

24. The communication network as described in Claim 23, wherein said front-end BTCP module parses a HTTP request from said client in order to determine which of said plurality of back-end web servers, a selected back-end web server, contains said data in order to process said HTTP request.

25. The communication network as described in Claim 23, wherein said front-end BTCP module extends said communication session to said selected back-end web server by handing-off an initial TCP state of said front-end BTCP module to a second BTCP module located at said selected back-end web server, said initial TCP state associated with a proper TCP state for said front-end BTCP module in said communication session, said front-end BTCP module further forwarding packets, including said HTTP request, from said

client after successfully handing-off said initial TCP state.

26. The communication network as described in Claim 25, wherein said second BTCP module understands said proper TCP state of said front-end BTCP module in said communication session and modifies headers in response packets from said selected back-end web server to reflect said proper TCP state.

27. The communication network as described in Claim 25, wherein said BIP module changes a destination address in forwarding said packets from said client.

28. The communication network as described in Claim 26, wherein said second BTCP module located at said selected back-end web server sends said response packets from said selected back-end web server to said client in a communication path that does not include said front-end node by changing headers of said response packets such that it appears the source of said response packets is said front-end node.

29. The communication network as described in Claim 22, wherein said content is partially replicated between each of said plurality of back-end web servers.